

Claims

What is claimed is:

1. A method for joining two value-ordered primary index tables T1 and T2 in response to a join command, T1 and T2 each having rows and T1 and T2 each having a first row, comprising the steps of:

- a. loading a subset of the rows of T1 into memory, where the subset is consecutive and includes the first row of T1;
- b. loading the first row of T2 into memory;
- c. comparing the loaded row of T2 to the loaded rows of T1;
- d. outputting a match, if any;
- e. recording a lack of a match, if any;
- f. loading a next consecutive row of T2 into memory;
- g. repeating steps c through f;
- h. loading new rows of T1 into memory, where the new rows are consecutive and are consecutive with the previously loaded rows of T1; and
- i. repeating steps c through h.

2. The method of claim 1, further comprising the step of:

- h'. discarding rows of T1 from memory, where the discarded rows are consecutive and include the first row that is in memory and wherein step i comprises repeating steps c through h and h'.

3. The method of claim 2, wherein the join command specifies one or more conditions on rows of T1 and discarded rows of T1 are stored if they meet the conditions.

4. The method of claim 1, further comprising the step of:

- f. discarding the previous row of T2 from memory and wherein step g comprises repeating steps c through f and f and step i comprises repeating steps c through h and f'.

5. The method of claim 1, wherein step h occurs after step c finds no matches when a predetermined number of matches have already occurred.

6. The method of claim 5, wherein the predetermined number of matches includes a first predetermined number of matches in a first portion of the loaded rows of T1 and a second predetermined number of matches in a second portion of the loaded rows of T1.

7. The method of claim 1, wherein the join command specifies one or more join columns for T1 and T2, further comprising the steps of:

j. halting the repetition of steps a-h when a predetermined number of comparisons do not result in a match;

k. sorting rows of T1 and T2 by their join columns, respectively; and

l. merge joining the sorted rows of T1 and T2.

8. The method of claim 1, wherein T1 and T2 are ordered by a date value.

9. The method of claim 1, wherein T1 and T2 each have one or more primary index columns and the join command specifies equality of the primary index columns as a condition of the join.

10. The method of claim 9, wherein the one or more primary index columns is a single column containing numeric values of 8 bytes or less.

11. The method of claim 1, wherein T1 has a unique primary index.

12. The method of claim 1, wherein T1 and T2 are portions of larger tables and are stored in a single storage facility.

13. The method of claim 1, further comprising the step of:

j. recording the maximum and minimum difference between the ordering columns of T1 and T2 for matching rows.

14. The method of claim 1, wherein the join command specifies one or more conditions on rows of T2 and recording a lack of match comprises recording a lack of match only when the loaded row of T2 meets the conditions.

15. The method of claim 1, wherein the join command specifies one or more join columns for T1 and T2, further comprising the steps of:

j. halting the repetition of steps a-h when a predetermined number of comparisons in a row do not result in a match;

5 k. sorting rows of T1 and T2 based at least in part on their join columns, respectively; and

l. merge joining the sorted rows of T1 and T2.

16. The method of claim 15, wherein the rows of T1 and T2 are sorted by a hash of their join columns.

10 17. The method of claim 1, wherein the join command specifies one or more conditions on rows of T1, rows of T2, and combined rows of T1 and T2, and outputting a match comprises outputting a combined row including one or more columns from the matched rows of T1 and T2 if each of the matched rows meets the conditions specified for rows from that table and the combined row meets the conditions for combined rows.

18. A database system for joining two value-ordered primary index tables T1 and T2 in response to a join command, T1 and T2 each having rows, the system comprising:

one or more nodes;

a plurality of CPUs, each of the one or more nodes providing access to one or more CPUs;

a plurality of virtual processes, each of the one or more CPUs providing access to one or more virtual processes;

each virtual process configured to manage data, including rows of tables T1 and T2, stored in one of a plurality of data-storage facilities, where T1 and T2 each have a first row in each data storage facility;

a reading join component configured to join rows from T1 and T2 by:

- a. loading a subset of the rows of T1 into memory, where the subset is consecutive for rows in that data-storage facility and includes the first row of T1;
- b. loading the first row of T2 into memory;
- c. comparing the loaded row of T2 to the loaded rows of T1;
- d. outputting a match, if any;
- e. recording a lack of a match, if any;
- f. loading a next consecutive row for rows in that data-storage facility of T2 into memory;
- g. repeating steps c through f;
- h. loading new rows of T1 into memory, where the new rows are consecutive for rows in that data-storage facility and are consecutive for rows in that data-storage facility with the previously loaded rows of T1; and
- i. repeating steps c through h.

19. The system of claim 18, wherein the reading join component is configured to join rows from T1 and T2 by:

- h'. discarding rows of T1 from memory, where the discarded rows are consecutive for rows in that data-storage facility and include the first row that is in memory and wherein step i comprises repeating steps c through h and h'.

20. The system of claim 19, wherein the join command specifies one or more conditions on rows of T1 and discarded rows of T1 are stored if they meet the conditions.

21. The system of claim 18, wherein the reading join component is configured to join rows from T1 and T2 by:

- 5 f. discarding the previous row of T2 from memory and wherein
 step g comprises repeating steps c through f and f' and step i comprises repeating steps c
 through h and f'.

22. The system of claim 18, wherein step h occurs after step c finds no matches when a predetermined number of matches have already occurred.

10 23. The system of claim 22, wherein the predetermined number of matches includes a first
 predetermined number of matches in a first portion of the loaded rows of T1 and a second
 predetermined number of matches in a second portion of the loaded rows of T1.

24. The system of claim 18, wherein the join command specifies one or more join columns for T1 and T2 and the reading join component is configured to join rows from T1 and T2 by:

- 15 j. halting the repetition of steps a-h when a predetermined number of comparisons do not
 result in a match;
 k. sorting rows of T1 and T2 by their join columns, respectively; and
 l. merge joining the sorted rows of T1 and T2.

25. The system of claim 18, wherein T1 and T2 are ordered by a date value.

20 26. The system of claim 18, wherein T1 and T2 each have one or more primary index columns and the join command specifies equality of the primary index columns as a condition of the join.

27. The system of claim 26, wherein the one or more primary index columns is a single column containing numeric values of 8 bytes or less.

28. The system of claim 18, wherein T1 has a unique primary index.

29. The system of claim 18, wherein the reading join component is configured to join rows from T1 and T2 by:

- j. recording the maximum and minimum difference between the ordering columns of T1 and T2 for matching rows.

5 30. The system of claim 18, wherein the join command specifies one or more conditions on rows of T2 and recording a lack of match comprises recording a lack of match only when the loaded row of T2 meets the conditions.

31. The system of claim 18, wherein the join command specifies one or more join columns for T1 and T2, and the reading join component is configured to join rows from T1 and T2 by:

- 10 j. halting the repetition of steps a-h when a predetermined number of comparisons in a row do not result in a match;
- k. sorting rows of T1 and T2 based at least in part on their join columns, respectively; and
- l. merge joining the sorted rows of T1 and T2.

15 32. The system of claim 31, wherein the rows of T1 and T2 are sorted by a hash of their join columns.

20 33. The system of claim 18, wherein the join command specifies one or more conditions on rows of T1, rows of T2, and combined rows of T1 and T2, and outputting a match comprises outputting a combined row including one or more columns from the matched rows of T1 and T2 if each of the matched rows meets the conditions specified for rows from that table and the combined row meets the conditions for combined rows.

34. A computer program, stored in a tangible medium, for joining two value-ordered primary index tables T1 and T2 in response to a join command, T1 and T2 each having rows and T1 and T2 each having a first row, the program comprising executable instructions that cause a computer to:

- a. load a subset of the rows of T1 into memory, where the subset is consecutive and includes the first row of T1;
- b. load the first row of T2 into memory;
- c. compare the loaded row of T2 to the loaded rows of T1;
- d. output a match, if any;
- e. record a lack of a match, if any;
- f. load a next consecutive row of T2 into memory;
- g. repeat steps c through f;
- h. load new rows of T1 into memory, where the new rows are consecutive and are consecutive with the previously loaded rows of T1; and
- i. repeat steps c through h.

35. The computer program of claim 34, further comprising executable instructions that cause a computer to:

- h'. discard rows of T1 from memory, where the discarded rows are consecutive and include the first row that is in memory and wherein step i comprises repeating steps c through h and h'.

36. The computer program of claim 35, wherein the join command specifies one or more conditions on rows of T1 and discarded rows of T1 are stored if they meet the conditions.

37. The computer program of claim 34, further comprising executable instructions that cause a computer to:

- f. discard the previous row of T2 from memory and wherein step g comprises repeating steps c through f and f' and step i comprises repeating steps c through h and f'.

38. The computer program of claim 34, wherein step h occurs after step c finds no matches when a predetermined number of matches have already occurred.

39. The computer program of claim 38, wherein the predetermined number of matches includes a first predetermined number of matches in a first portion of the loaded rows of T1 and a second predetermined number of matches in a second portion of the loaded rows of T1.

40. The computer program of claim 34, wherein the join command specifies one or more join columns for T1 and T2, further comprising executable instructions that cause a computer to:

- j. halt the repetition of steps a-h when a predetermined number of comparisons do not result in a match;
- k. sort rows of T1 and T2 by their join columns, respectively; and
- l. merge join the sorted rows of T1 and T2.

41. The computer program of claim 34, wherein T1 and T2 are ordered by a date value.

42. The computer program of claim 34, wherein T1 and T2 each have one or more primary index columns and the join command specifies equality of the primary index columns as a condition of the join.

43. The computer program of claim 42, wherein the one or more primary index columns is a single column containing numeric values of 8 bytes or less.

44. The computer program of claim 34, wherein T1 has a unique primary index.

45. The computer program of claim 34, wherein T1 and T2 are portions of larger tables and are stored in a single storage facility.

46. The computer program of claim 34, further comprising executable instructions that cause a computer to:

- j. record the maximum and minimum difference between the ordering columns of T1 and T2 for matching rows.

47. The computer program of claim 34, wherein the join command specifies one or more conditions on rows of T2 and causing a computer to record a lack of match comprises causing it to record a lack of match only when the loaded row of T2 meets the conditions.

48. The computer program of claim 34, wherein the join command specifies one or more join columns for T1 and T2, further comprising executable instructions that cause a computer to:

j. halt the repetition of steps a-h when a predetermined number of comparisons in a row do not result in a match;

5 k. sort rows of T1 and T2 based at least in part on their join columns, respectively; and

l. merge join the sorted rows of T1 and T2.

49. The computer program of claim 48, wherein the rows of T1 and T2 are sorted by a hash of their join columns.

50. The computer program of claim 34, wherein the join command specifies one or more conditions
10 on rows of T1, rows of T2, and combined rows of T1 and T2, and causing a computer to output a match comprises causing a computer to output a combined row including one or more columns from the matched rows of T1 and T2 if each of the matched rows meets the conditions specified for rows from that table and the combined row meets the conditions for combined rows.